Good practice in information handling:
Audit logging and incident handling

For staff and contractors tasked with implementing data security

This document is one of a series of good practice guides to help schools, colleges and universities protect personal and sensitive data. Building on good practice from industry and central government these guides describe procedures and possible technical and operational solutions that can help organisations reduce the risks of data security incidents and comply with current legislation.

Produced by Becta on behalf of the Department for Children, Schools and Families, these guides have been reviewed and updated with feedback from a number of cross-sector organisations including DCSF, DIUS, JISC Legal, The Information Authority and JANET(UK), as well as from schools, local authorities, RBCs and suppliers.

For further information on these guides, please see http://www.becta.org.uk/schools/datasecurity and http://www.becta.org.uk/feandskills/datasecurity

Contents

Key points

This guide outlines why organisations need to have a policy and procedures for audit and event logging. Educational organisations must keep audit logs to help detect and respond to security incidents, and to provide evidence of accidental or deliberate security breaches, for example, loss of personal data or breach of an acceptable-use policy.

Adopting ISO 27001/27002 [http://en.wikipedia.org/wiki/ISO/IEC_27000-series] will help make an organisation compliant with the Government's current guidelines on data security. We acknowledge, however, that to follow these standards may be difficult, expensive and time-consuming for many institutions. This document, therefore, provides good practice guidance to assist staff and contractors tasked with implementing data security and putting in place a system of audit logging and incident handling. It contains:

- help with planning and implementing a basic audit logging infrastructure
- information on security event auditing – collecting ICT system events and reviewing their impact
- help with devising a monitoring strategy
- details of the actions that you need your logging system to detect
- good practice in audit logging
- help with planning ahead, so you can respond effectively to a breach of data security.

1 The need for audit logging

Collecting audit logs is crucial to providing a safe and secure ICT infrastructure for educational environments. Educational organisations need to put in place a system that consolidates all of the log data recorded by their information management systems, learning platforms, portals and various hardware devices.

1.1 Which devices do we audit?

A typical network arrangement for an educational organisation has a number of items of hardware that audit logs should be collected from, including:

- hardware and software-based firewalls
- web servers
- central/domain controllers
- authentication servers
- management information system servers
- learning platform web servers
- web portal, database, query and index servers
- mail servers and web mail servers

- file servers
- routers
- DHCP servers
- Network Address Translation (NAT) devices
- networked PCs and other connected devices
- audit servers.

Logging infrastructures must therefore be able to operate in a variety of distributed networks and collect secure and evidential quality logs. An organisation's network or managed service provider may already carry out some logging.

Logging produces large amounts of data. Organisations do not need to keep it forever (and, in most cases, must not in order to comply with the Data Protection Act 1998). Organisations should only hold logs for the length of time stated in their audit policy. The section on audit logging good practice later in this document discusses the things organisations should consider. The Records Management Society has written a useful guide for schools, *Records Management Toolkit for Schools* [http://www.rms-gb.org.uk/resources/848], which includes a section on retention.

2 Planning an audit-logging infrastructure

The information below is based on industry good practice and is intended as a guide. Organisations need to define their own policies and procedures to meet their own requirements.

There are three ways to monitor systems for security breaches:

- Network level TCP/IP
- Server and application
- Process-specific.

2.1 Example of an audit logging infrastructure

Currently, many educational organisations do not have an audit logging policy or consolidated auditing infrastructure. They may simply log the firewalls supporting the internet and email.

Figure 1 shows the infrastructure that organisations will need to implement to provide audit and event logging. In this infrastructure, organisations will need to manually consolidate log data.

A critical aspect of any audit/log collection system is accurate time stamps, so network time synchronisation is essential.

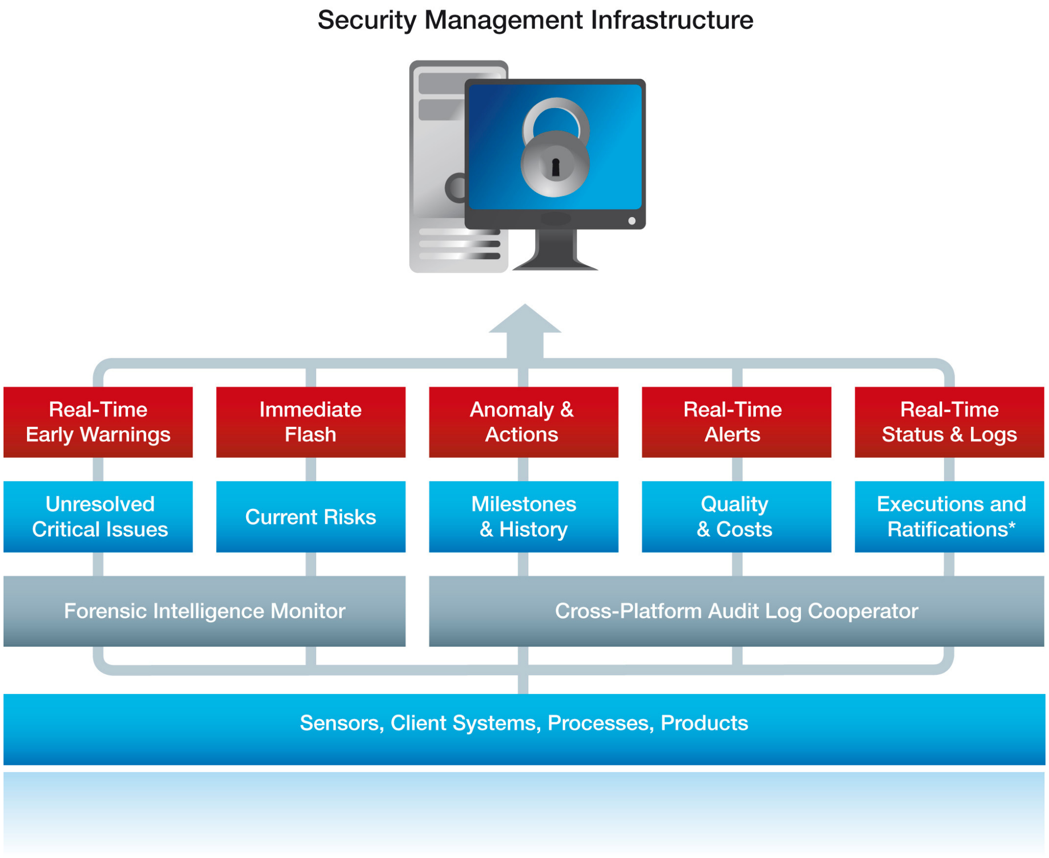
Security Management Infrastructure
Real-Time Early Warnings
Immediate Flash
Anomaly & Actions
Real-Time Alerts
Real-Time Status & Logs
Unresolved Critical Issues
Current Risks
Milestones & History
Quality & Costs
Executions and Ratifications*
Forensic Intelligence Monitor
Cross-Platform Audit Log Cooperator
Sensors, Client Systems, Processes, Products

Figure 1: An audit and event logging infrastructure

The infrastructure shown in Figure 1 may require organisations to install extra hardware and software to provide evidential quality storage and archiving.

2.2 Implementing an audit logging infrastructure

First steps to implement an infrastructure include:

- listing critical systems (including those with personal data) and determining what logging is turned on, where this log data is stored, how long it needs to be kept, the format, who owns the system and who can access it
- calculating the amount of data produced to work out network bandwidth and storage space requirements and recording format
- getting hold of the necessary servers, hubs, network-attached storage and firewalls to build a secure internal area for these items. Organisations may need to build more than one audit/logging area due to the distributed nature of their infrastructure.
- naming staff who have responsibility for operating the infrastructure, including the information that is to be reported, archiving processes, and procedures for resolving discoveries and remediation requirements.

2.3 Audit and logging tools

Raw log data is hard to analyse and interpret; manually consolidating logs is also hard and different platforms use different formats. In order to find security exploits, staff need advanced ICT and statistical skills. Organisations may need to provide authorised staff with additional tools to carry out cross-platform query and reporting. A number of software products are available to help collect and analyse audit and log data and examples of such products are provided here as an aid:

- Advanced Log Analyser [http://www.abacre.com/ala]
- CA Spectrum [http://www.ca.com/gb/products/product.aspx?id=7832]
- Enterasys Dragon by Enterasys [http://www.enterasys.com/products/advanced%2Dsecurity%2Dapps]
- Event Analyst [http://www.doriansoft.com/totalsolution/index.htm]
- Event Log Explorer [http://www.eventlogxp.com]
- Exchange Log Analyzer [http://www.mechanicalminds.com/site/ela]
- Snare (Open Source) [http://www.intersectalliance.com]
- StealthWatch [http://www.lancope.com]
- WizRule [http://www.wizsoft.com].

Note: Becta has not conducted a formal evaluation of these products and, therefore, does not recommend any specific solution. Other suitable products are available that are not listed in this document.

Some of these tools need extensive technical knowledge of the target system, however; others can talk to the target system, perform the necessary queries and provide an integrated report. This further reduces the need for staff to have technical knowledge, training and forensic ICT experience.

2.4 Barriers to implementing an audit logging infrastructure

The main barriers to implementing an audit log infrastructure are the time needed, concerns from staff about logging and access to log data, compliance with the Data Protection Act 1998, and the extra storage space needed to hold logs. Organisations need to consider all these issues before implementing audit logging.

Implementation projects have often failed for the following reasons:

- Lack of support from senior management and lack of directives mandating that all involved parties actively support the tasks
- Lack of management enforcement of the adopted policies
- Issues relating to data ownership and network bandwidth availability
- Inadequate funding to upgrade servers owned by others
- Lack of skilled staff and the ability to dedicate the required time, in spite of this being of the highest priority
- Lack of user engagement.

3 Security event auditing

An organisation’s network or managed service provider may already carry out some security event auditing, so organisations should check what they already have in place.

3.1 What is security event auditing?

Security event auditing is the process of collecting ICT system events and reviewing the impact of these against security policy. Context is gained by linking the series of events with users’ actions.

The process of auditing determines compliance or non-compliance to established security policies and procedures. Regular audits are recommended as part of running an efficient and well-controlled ICT operation.

3.2 Identifying vulnerabilities

Reviewing the security configuration will indicate whether it is set in accordance with a defined baseline guided by the security policy. Aspects of the security configuration include security level, audit policy, password characteristics, registry, file/directory permissions, user accounts, groups, rights, privileges and network configuration.

3.3 Identifying suspicious activities, break-in attempts and security breaches

The purpose of the audit is to identify any events that indicate suspicious activity. These would include, for example, users who repeatedly failed to log on, users who logged on at unusual times or logged on from unknown remote systems, users who failed to open files or folders due to insufficient permissions, unusual use of admin privileges, and users who have repeatedly attempted to access system services, but failed due to insufficient privileges.

Most of these events are likely to be the result of a normal system activity. Mistyping of passwords or accidentally trying to open files without appropriate permission are common mistakes. A good indicator of an attempted break-in is the number of repeat attempts and organisations should know enough about their own typical routine activity to enable them to determine whether any events warrant deeper investigation. It is therefore important to periodically scan audit logs, even when no alert has been raised.

Organisations should also look for possible intrusions from the outside, and find out whether an internal user has performed an unauthorised activity that violates the security policy. In other words, checking whether the security of a system has been breached and, if it has, determining exactly what has happened and where the first point of breach occurred.

To be able to do this, the computer system must have an event-logging facility to record the occurrence of significant events in the first place. The more sophisticated the event logging, the sooner you will detect an unauthorised activity.

3.4 What are the main audit data sources?

3.4.1 Security configuration snapshots

In order to identify system vulnerabilities, organisations need to collect and review relevant security information, including security level, audit policy, password characteristics, registry, file/directory permissions, user accounts, groups, rights, privileges and network configuration. This is the first type of audit data source, often referred to as 'security configuration snapshots' when captured and stored with a time stamp.

3.4.2 Event logs

In order to identify suspicious activities, break-in attempts and security breaches, organisations will need to collect and review all the significant events recorded by their event-logging facility. Note that the auditing parameters need to be appropriately configured in the first place, so that the event-logging facility records all the significant events you want to monitor. This is the second type of audit data source, often referred to as event logs. On Windows, these are known as system, application or event logs, and syslogs on Linux, Mac and UNIX.

3.5 What is event logging?

Event logging is the process of noting the occurrence of a significant event and recording it in a persistent medium. Each event is recorded in a log. Each record written to the log is referenced by date and time.

3.5.1 What kinds of events are logged?

In a security event log, you are likely to see security-relevant actions such as:

- users logging on and off the system
- changes made to system security and user privileges
- attempts to access file, directories, printers and other system objects that are under audit control.

These types of events inform ICT network managers and/or dedicated security managers how users and processes are attempting to access and use the system. The security log therefore contains a persistent record of 'who is doing what, when and where from'. Such an audit trail would give an indication of repeated attempts to illegally log on to a remote computer, gain access to secured files or install unauthorised software. On a more mundane level, it would simply point to someone who seems to be printing out a lot of unnecessary documents, or spending too much time on the internet.

The auditing policy determines which events are recorded in a log. Specific system and user events are pre-selected by the security administrator based on how the system is configured and for what it is used.

3.6 Why use event logging?

It is impossible to guarantee that any computer system is 100 per cent secure. There will always be security flaws that can be exploited, and unfortunately the greatest risks to security are the users themselves.

If unauthorised entry and access to protected data cannot be prevented, such problems at least need to be recorded and tracked in an event log for the purpose of revealing the security flaws in the systems and possibly identifying those (humans or computers) that have exploited these flaws.

To summarise, event logs provide the information required to identify attempted attacks, to investigate what happened when an incident has occurred, and to potentially provide evidence in support of an investigation. They may also identify areas where a system is not performing as users expect and where the system should be fixed before users develop their own insecure workarounds.

3.7 International standards for event logging

The best practice standard for measuring and performing event logging is based upon the Common Criteria for Information Technology Security Evaluation (ISO 15408) [http://www.cesg.gov.uk/products_services/iacs/cc_and_itsec/cc_levels.shtml]. Education systems will generally require a minimum security level based on the classes defined in this. The Common Criteria security evaluation class requires that a computing system must have an auditing mechanism with the following minimum capabilities:

- The system has the ability to record all security-related events that occur on the system in the form of audit records
- The system provides a way for the audit records to be reviewed by the system administrators
- The auditing software and logs must be protected by the operating system from unauthorised access and modification, and access must be limited to authorised system administrators
- A mechanism must exist that allows the selection of security events to be audited
- The system must be able to audit individual users.

The Common Criteria stipulate that each audit log record must contain the following information:

- Date and time of the event's occurrence
- Unique ID of the user creating the event
- Type of event
- Success or failure of the event
- Origin of the event (user, system, terminal)
- Name of the object accessed (a system file, piece of data or computing process)
- Description of any modifications made to security databases.

For these requirements to be met, the logging system must be able to monitor both its own activities and those of all local and remote users connected to the system. It must also be able to report events pre-selected by the system administrators to one or more central structures. The logs and the logging mechanism must be guarded using the highest sensitivity level possible for system objects.

On protected systems, it is possible to track a user's access of files and directories, printers, network volumes and shares, and attempts to modify any security aspects of the system such as changing file permissions, adding a user account or privilege, or disabling auditing. All such actions may be recorded as events in a security audit log.

3.8 The limitations of event logging

All of this information does not necessarily indicate the specific events to log, or how to classify the events. For example, is a user log-on and log-off a single event with two states, or two separate events? The audit events recognised by a system depend entirely on the capabilities and sophistication of the system's components and security mechanisms.

Changes to the security configuration should be recorded in the event logs, however, inherent weaknesses are observed in most event-logging facilities. It may not be possible to record all user and process actions in sufficient detail. Because of the deficiencies in the event logs themselves, it becomes necessary to rely additionally on security configuration snapshots that identify vulnerabilities.

The following examples may help to illustrate the problem.

3.8.1 Example 1

If a security administrator changes the value for minimum password length (as part of the password policy settings), the security event log will record that a policy change has occurred, but not record what exact change was made. To be able to detect such a change would require a snapshot of the policy settings before and after the change was made.

3.8.2 Example 2

The default logging mechanisms capture commands that were executed but not always the parameters associated with those commands. If the access permission on sensitive files was changed temporarily for a few hours during the working day to gain inadvertent access to secured files, and if the snapshot of the file system is only collected at the end of each working day, the fact that permission on those files was changed and reset to its original value after a few hours will not be detected.

4 Monitoring strategy

Organisations must tell users that they are being monitored. This is a requirement of the Data Protection Act 1998. Organisations should also have acceptable-use policies (AUPs) in place. Monitoring provides a significant deterrent effect. Users who know that their actions may be being monitored are less likely to breach regulations. Organisations may see significant changes in user behaviour after the introduction of a monitoring system. It is not necessary to monitor every computer in the network to provide this effect.

Where users are not deterred, monitoring also provides support for corrective action, as the data gathered is crucial evidence. Evidence of unauthorised activity produced by the monitoring system is usually sufficient to take administrative action against an offender. This level of response, together with an awareness campaign about the monitoring, is usually enough to make the most stubborn user (or administrator) stop unauthorised activity.

4.1 Detection

Detection in a broad context means the identification of activities of note. Noteworthy activities may not necessarily be considered misuse or an intrusion. For example, detecting accesses to a mission-critical file may be misuse if the user account accessing the file is unauthorised to do so. However, simply counting the number of unique individuals that access a mission critical file can indicate that critical data is too widely available within a given organisation, local authority or other institution. An example would be the number of individuals accessing protected data and attempting to download the entire database. This is a noteworthy activity and may be tracked by a monitoring system.

4.2 Response and notification

Some monitoring systems are able to react after detecting misuse. These reactions may be automated or manual and include both local and remote actions and notifications. The notifications or alerts are similar to network-based monitoring systems and usually include:

- pager
- SMS

- SNMP Trap
- on-screen
- audible
- email.

Once they have detected misuse, most commercial products include the following response actions:

- log off user
- shut down system
- disable account
- execute local script.

Organisations setting up systems to perform automated actions should remember that these systems sometimes wrongly identify cases of misuse. This can result in systems being disabled without need or in administrators being notified out of hours unnecessarily.

4.3 Damage assessment

If an organisation has correctly identified a security incident or data loss, the first question the organisation should answer is: ‘What was the extent of the damage?’

Typically, organisations must pull files from storage and may spend hundreds of staff hours looking through data to assess the extent of the damage. Staff must spend additional hours analysing databases and file systems searching for unauthorised changes. If the damage was unauthorised disclosure, only archived event log data will be of any value. If an organisation has no log archive, then it will not be able to determine the extent of the damage.

A key part of monitoring is maintaining an archive of information that can be queried using data forensics tools. An event log archive helps answer the following questions:

- Which computers were accessed?
- Which sensitive files were accessed and/or modified?
- What methods and tools were used to gain access?
- Was the individual working alone?
- How long have the events been going on?

4.4 Event anticipation

Many events are characterised by a set of preliminary activities before the actual loss is incurred. For example, a user who is looking for sensitive data with malicious intent may start a systematic search of systems before finding critical data. If this is spotted, the user can be locked out before the data is compromised.

Many breaches will, however, be accidental. An example of this is when a member of staff downloads data from the management information system. Frequently, they take this data with them by saving it on an encrypted USB portable drive. This action is routine for one or two records. However, doing this for the entire database should be an exception and require authorisation by the Information Asset Owner. This type of activity would be detectable by the monitoring system. Once the detection is made, the user can be warned while still in the process of generating the report or creating the file.

4.5 Corrective resolution support

The term ‘corrective resolution’ refers to action(s) that may be taken against individuals as the consequence of their culpability in the loss of data, a breach of security or a breach of security policy. Corrective resolution may cover exclusion, suspension, restriction of privilege, restriction of use, termination of contract, prosecution or any other measure required to be enacted. Severity of corrective resolution must be determined based on the context of the incident and within the judgement of the relevant authority.

Server tools can provide data to support corrective resolution, that is, action which follows detection of an infringement by a specific individual. This data includes access patterns to files and computers with specific dates and times. This data may not be sufficient by itself to lead to disciplinary action but, in conjunction with other evidence, it can indicate specific computer activities.

Note, however, that there are rules that control the admissibility of monitoring data as evidence in a court of law. Forensic evidence such as logs must pass several tests including chain of custody and integrity tests. Organisations must be able to reasonably protect the data from modification and account for its location and who had access to it between the time it is collected and the time it is submitted in court as evidence. Server and application-based monitoring can provide the secure collection and storage mechanisms so that the data may be admitted as evidence in court.

You may need to employ the assistance of a forensic analyst to explain the actual log events and the relationship between system events. A critical aspect of any audit/log collection system is accurate time stamps of these events. Network time synchronisation is a priority, or you will need to calculate the time difference for each event across the various systems. Without consistent time stamps (such as network time domain), it is almost impossible to provide the required assurances for court evidence.

5 Defining requirements for security incident handling

5.1 Goals for monitoring with respect to incident response preparation

It is a legal requirement of the Data Protection Act 1998 to have a statement of intent that tells staff what kinds of actions are monitored and the level of detail involved.

Monitoring of this sort should be made clear in any staff handbook, and in any acceptable-use policy, together with the response that the organisation will take in cases of misuse.

5.2 Detection requirements

It is essential to specify detection requirements. Organisations should decide what they want to obtain from data so that they can gather data on the right events.

Organisations should use their statement of intent to help define the actions that they would like to detect, starting at a high level and moving into more detailed requirements as the process develops.

5.2.1 Insider threat detection requirements

Insider threat detection requirements cover the actions of authenticated users. Trend and behavioural analysis is often used to detect insider misuse. Investigators can prevent damage by detecting suspicious behaviour in access patterns to critical data and systems.

Typical monitoring requirements of insider threat detection systems are:

- tracking of access patterns to critical data (that is, protected data) and systems
- detection of common vulnerabilities used to gain privileges as an authenticated user.

(Note: Unauthorised use of privileges is a general class of detection requirement serviced by server and application-based detection mechanisms.)

5.3 Compliance monitoring requirements

Compliance monitoring helps to make sure that people and processes are following security policies. Policies are there to provide a verifiable level of protection in a network. Compliance monitoring can be provided through behavioural monitoring or static configuration analysis.

To establish requirements for compliance monitoring, organisations should consider which security policies, if ignored, would result in the most significant losses.

Any monitoring system should also detect when people or systems are not following the rules set out in the audit/logging policy.

5.4 Incident response requirements

Response requirements often start out being very ambitious (such as automatically logging out users and shutting down access to systems) but as the security co-ordinators begin to understand the risks associated with automated responses, they can relax these requirements significantly.

Most response requirements should be associated with escalation procedures designed to make the system most effective. Once again, these requirements should relate to your statement of intent.

If your primary reason for using logging is to record breaches, then your response requirements may focus on escalating administrative actions as an investigation proceeds. If your primary reason is real-time analysis, then you need to carefully consider the timing requirements for manual response and balance automated responses against the risks they pose.

Requirement: When a user falls under suspicion during the course of an investigation, the user’s account shall be disabled on the approval of the ICT network manager or site security manager. The account shall remain disabled until authorised for re-enablement by the appropriate party.

(Note: This requirement relates to escalation procedures and prevents continued misuse during an investigation. However, there may be circumstances where the goal is to not alert the individual under review and allow continued access for surveillance, so this requirement could be restated. This will be dependent on the specific circumstances involved.)

5.5 Resource classification

Not everything can be monitored, so you want to be able to focus analysis on critical assets and systems containing sensitive data. You will need to differentiate critical data from non-critical data, and sensitive data from non-sensitive data. The process of defining data is known as resource classification. Definitions must be established for identifying, controlling and handling the different classifications. These include information elements such as resource, assigns, grants, requirements and recovery. The information populated in these categories establishes the requirements for each resource in each classification level. These are defined as:

- **Resource –** This is the definition of each resource. Define these as necessary to match the critical assets in your organisation. For example, all HR data files, protected directories, website content files, and so on.
- **Assigns –** This is the person who assigns the classification to the resource, for example, the Senior Information Risk Owner (SIRO), the

department manager, and so on. It is important to control data classification carefully and commonly between all users who share that data in different organisations. Data controls cost time and money so you need to be able to balance security with access requirements.

- **Grants –** This is the person who grants access to the resource (for protected data this is usually the Information Asset Owner (IAO)). This is a very different role from the assignment role. The assignment role requires decision-making abilities and a higher-level view while the granting role is more about following instructions. The significance of the granting role depends on the exact nature of the control requirements.
- **Requirements –** These are the control requirements. They cover as broad a range as necessary to protect the data at its classification including who can access the data and how exceptions are handled. For example, the restriction 'Only persons in the organisation's leadership team may open and read summary information management data' could allow as an exception 'Staff may access personally identifying data for individuals in their organisation with permission from the granting authority'.
- **Recovery –** If data is important enough to be controlled, then it should have recovery and back-up requirements.

5.6 Platform coverage requirements

Platform coverage requirements describe the functions and systems that organisations want to monitor. These include network protocols, operating systems, computers and applications.

Analysis tools will not be compatible with all network hardware and software. Organisations should identify what they need to monitor, establish operational requirements, prioritise the critical hardware or software most at risk, and then select appropriate analysis tools.

Typical requirements for platform coverage are that monitoring is:

- required on different platforms (Windows, Linux, Mac OS)
- deployed to sensitive applications first, and then migrated to support infrastructures and user applications.

(Note: This requirement sets a priority order for the deployment and relative importance of the various systems.)

5.7 Audit source requirements

Once you have defined your platforms, you need to define which audit sources you need to monitor on those platforms. Many servers have multiple operating system and application logs. The different logs have different levels of information and security associated with them. This requirement will enable you to enumerate the different audit sources that your monitoring procedures will need to cover.

Requirement: The logging software shall have the ability to monitor the various logs from operating systems, routers, wireless access points, firewalls and applications (management information system and learning platform).

5.8 Corrective resolution requirements

As you will be using data collected during monitoring to enforce potential action against individuals, you need to take care to protect this as evidence. This includes selecting appropriate audit sources and tools that protect data sufficiently for admissibility in court. Many of the requirements for corrective resolution are process-related and these requirements establish the level of information that needs to be available at various points during the investigation, or for reference by the relevant authorities.

Requirement: When a user falls under suspicion during the course of an investigation, all raw data relating to the investigation shall be stored for evidence.

(Note: Requirements for evidence place constraints on the quantity and quality of data as well as how it is handled. For further information, refer to the Association of Chief Police Officers' *Good Practice Guide for Computer-based Electronic Evidence* [http://is.gd/nS8s].)

6 Good practice for audit logging

As we have seen, audit logging can be complex and policies will differ between organisations as infrastructures differ and a balance is made between existing resources and risk management. In this section we suggest some good practice for audit logging that should be considered and implemented where relevant.

6.1 Prerequisites

An organisation's monitoring strategy should establish a security configuration baseline. Organisations should collect security configuration snapshots and event logs. These snapshots can be used to determine whether the underlying system configuration has been changed. Organisations can also use snapshots during security audits.

An organisation's monitoring strategy should also determine what events they need to monitor and how often they collect event logs. Organisations may also use sophisticated tools and techniques to review certain events online in real-time.

6.2 Archiving strategies

Organisations can chose from several archiving strategies. Each strategy is based on two elements: archiving and centralising.

Archiving is copying logs to another location and clearing them. Centralising involves moving logs to a central archive to free disk space.

6.2.1 Table showing some strategies and their advantages and disadvantages

Strategy	Advantages	Disadvantages
Archive once each hour Centralise once each day	Small raw log files, so easy to analyse Not much disk space needed for live log files	More files to archive On local machine for longer so attacker could modify or delete them before they are centralised
Archive once each hour Centralise immediately	Small raw log files, so easy to analyse Not much disk space needed for live log files Archived log files exposed to attack for less time	Immediate centralisation will spike network traffic each hour
Archive once each day Centralise once each week	Small raw log files, so easy to analyse Not much disk space needed for live log files Easy on resources	Archived log files at risk of attack for long periods

To select a strategy, organisations need to balance the availability of resource against the requirements to collect logs in a timely and secure fashion.

Organisations should retain logs as specified in their audit/log policy. The following is good practice for log management:

- An organisation’s audit/log policy should state which systems are logged and the retention periods for each system.
- If an organisation collects the audit logs from their systems for a month and during that month no incidents have occurred, they could archive the data offline and retain it for one academic year (recorded in one-month intervals). If an event happened in a subsequent period, it would be necessary to go back to check if any previous pattern existed or if this was an isolated event. In such cases the archive schedule would likely provide an acceptable and relevant data set. This should be outlined in the audit logging policy.

- If a breach of the acceptable use policy had occurred, organisations would need to extract all the related data surrounding that particular incident and create a case file. Organisations should keep this file for the length of time specified in their retention policy.

6.3 Rolling over the archive from online to offline storage

To carry out an investigation organisations will need to be able to access the data they have archived. Some of this data will be kept online for easy access and some will be stored on offline media. Again, these requirements will depend on the organisation or recent incidents.

A very common practice is to keep 30, 60 or 90 days online. Then, every 30 days or so, depending on requirements, put two copies of the last 30 days onto separate offline storage media and store them in separate secure locations.

6.3.1 Protecting data integrity

The integrity of data integrity is dependent on how quickly organisations copy it and the mechanisms and processes they use.

Offline data should be copied onto secure Write Once Read Many (WORM) media using an optical drive. This ensures that data remains unchanged.

Organisations must be able to show that data was created by a secure data source, copied very quickly to a secure central server, and then put on permanent media so that the data cannot be subsequently changed.

6.4 Disk space

Organisations will need to set aside disk space for both local collection on the target machine and for archiving. Organisations can estimate the amount of space the target system needs by calculating the amount of data it creates between each centralisation run (as defined by the archiving strategy). For example, if organisations archive logs once an hour and centralise once a day, they need to estimate how much data might be created in a day before the logs are centralised.

Organisations should account for additional risks when determining how much disk space is required for logs. For example, if a server is suddenly unable to send data, it must be able to store enough data until communication is restored and data is resynchronised, or until a process can be put in place to protect the data manually.

To work out the online disk space required, multiply the number of log files by the average daily log size and the duration of the storage requirement:

Disk space = (number of log files) x (average daily log size) x (number of days)

If organisations do not have enough disk space available on the target machines, disk space will run out and may cause your machine to stop and possibly shut down.

6.5 Audit data management

Audit data management is the practice of securely collecting and archiving audit (event) data for the purposes of analysis, investigation and corrective resolution. Data management also addresses the issue of collecting audit data in a way that does not affect the operation of your devices and network. It is important that the data is collected securely so that it can withstand cross-examination if it is ever necessary for a case to be taken to court.

Data should be gathered in a secure manner that is tamper-resistant. It is critical that the analysis methods are incapable of changing or altering the integrity of the data.

An effective audit policy is one that gathers just the right amount of data, such that you do not gather so much data that you suffer performance problems, and not so little data that events cannot be proved beyond doubt.

7 Building an effective security incident response capability

Organisations should establish a security incident response capability and, depending on the severity of the security incident, it may also take the form of disaster recovery. The size of an organisation and the specific skills of its staff may necessitate engagement with suppliers, the local authority or a network provider for an effective security incident response.

A prerequisite for an effective security incident response is the detection of the security incident in the first place. Thereafter, the effectiveness of the response team should be measured based on the extent of damage that resulted from each incident. The sooner the incident is contained, the lower the risk of harm to individuals or the organisation through financial or reputation loss or data compromise.

Good practice highlights the following components for the successful resolution of an incident:

- Management commitment, in terms of human resources, budget and priority
- A resolution team of technical and legal experts
- A person who is primarily responsible for each incident
- A communications plan, including escalation procedures and interfaces with inter-departmental and law enforcement agencies
- Plan of action for rapid resolution
- Plan of action for non-recurrence

- Knowledge base of past security incidents, including steps taken for resolution and non-recurrence
- An awareness campaign.

It is recognised that not all resources may be available and further guidance and support may be necessary to facilitate all the components outlined above.

7.1 Management commitment

Within the educational environment, strong emphasis is required from all responsible individuals in managing and responding to potential threats and risks to ICT. Senior representatives from all departments should establish a formal steering group to analyse the current system threats. It should cover the following functions and activities at a high level: data management, infrastructure, networks, systems and database administration, applications development, security, audit, legal and compliance.

7.2 The resolution team

The steering group should define the composition of the team that will address the technical, process and legal aspects of the issues at hand and expedite the solution.

7.3 Communications plan

Any activity as part of incident response is sensitive in nature. Improper publicity may harm the reputation of the institution. What is needed is a formal plan for when a security incident occurs.

The plan can be constructed by envisaging, provisioning and rehearsing for the events that occur when a security incident arises. Personnel contact lists, escalation procedures, operational responsibilities and guidelines should be formulated in advance and rigorously adhered to. Changes in plan should be avoided unless the situation leaves no other alternative.

Underpinning all of the above is a process of communicating your procedures to all relevant members of staff in the organisation, associated parties and law enforcement agencies to ensure:

- rapid resolution
- non-recurrence
- limitation of damage
- protection of all users.

7.4 Documentation

All actions taken during security incidents should be recorded and archived securely to ensure that complete evidential quality is maintained. Coupled with the recorded evidence of the event, a debriefing operation should be conducted to evaluate the performance of the team and effectiveness of resolution(s).

Subsequent to the incident, debrief material should be analysed to ascertain whether changes in infrastructures, operational practices or policies are required. Suggested improvements may then be fed back into the processes and procedures to enhance the response and reduce the chances of it happening again.

7.5 Awareness campaign

An awareness campaign must be developed that identifies the security incident team, contact points or agencies required to be informed on discovery of a suspected event. This should be disseminated to any party that may have cause to use such a system. This may include (but not be limited to) organisation staff, support contractors, local authority employees, parents, governors and (where appropriate) learners.